Create account Log in

Article Talk Read Edit View history Search

WIKIPEDIA
The Free Encyclopedia

Navigation
Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia
Wikimedia Shop

Interaction
Help
About Wikipedia
Community portal
Recent changes
Contact page

Tools
What links here
Related changes
Upload file
Special pages
Permanent link
Page information
Data item
Cite this page

Print/export
Create a book
Download as PDF
Printable version

Named data networking

From Wikipedia, the free encyclopedia

This article has multiple issues. Please help **improve it** or discuss these issues on the **talk page**.

- This article **is written like a a personal reflection or opinion essay that states the Wikipedia editor's particular feelings about a topic, rather than the opinions of experts**. *(March 2014)*
- This article **may be too technical for most readers to understand**. *(March 2014)*
- This article **does not cite any references or sources**. *(March 2014)*

Named data networking (also content-centric networking, content-based networking, data-oriented networking or information-centric networking) is an alternative approach to the architecture of computer networks. Its founding principle is that a communication network should allow a user to focus on the data he or she needs, rather than having to reference a specific, physical location where that data is to be retrieved from. This stems from the fact that the vast majority of current Internet usage (a "high 90% level of traffic") consists of data being disseminated from a source to a number of users.

The contemporary Internet architecture revolves around a host-based conversation model, created in the 1970s to allow geographically distributed users to use a few big, immobile computers. The content-centric networking seeks to adapt the network architecture to current network usage patterns.

Content-centric networking comes with potential for a wide range of benefits such as content caching to reduce congestion and improve delivery speed, simpler configuration of network devices, and building security into the network at the data level. However, the change of communication paradigm may pose problems for certain types of network activities, for instance for real-time multimedia applications, but recent research indicates these applications are feasible. Furthermore, building content routers that support content-centric networking at high speed is still an open problem to solve that has gained research interest only recently.

Application-layer designs have also been proposed for deploying a content-centric interface. This has benefits such as easier deployment, backwards compatibility and more flexible delivery support. The Juno middleware offers applications a content-centric request/reply interface that can be utilised alongside existing content providers. Juno also introduces the concept of a delivery-centric interface, which extends the traditional content-centric interface to allow applications to stipulate multiple diverse delivery requirements that place certain constraints on how the content should be provided. For instance, these constraints can deal with such things as performance, resilience, security, monetary cost and anonymity. Through such an interface, applications can shape how the underlying delivery is performed without needing to handle such concerns themselves.

History [edit]

The philosophy behind content-centric networks was pioneered by Ted Nelson in 1979 and later by Brent Baccala in 2002. In 1999, the TRIAD project at Stanford proposed avoiding DNS lookups by using the name of an object to route towards a close replica of it. In 2006, the DONA project at UC Berkeley and ICSI proposed a content-centric network architecture, which improved TRIAD by incorporating security (authenticity) and persistence as first-class primitives in the architecture. In 2009, PARC announced their content-centric architecture within the CCNx project, which is led by Van Jacobson, a research fellow at PARC. On September 21, 2009, PARC published the specifications for interoperability and released an initial open source implementation (under GPL) of the Content Centric Networking research project on the Project CCNx site.

How it works [edit]

The existing Internet is a tree of physical equipment to connect streams of packets from any leaf, to any leaf. It is efficient for communication, but not for distribution. The general proposal of content-centric networking recognizes that a great deal of information is produced once, then copied many times. Therefore, it makes sense to distribute the copying and any related activities into the network's tree of equipment. In many cases, substantial storage is already available, and could be used more efficiently if it could recognize particular content and only keep one copy of it. Since the network equipment is tree-shaped, it naturally scales content delivery to the size of the audience, and simultaneously reduces up-stream equipment to just the minimum needed to produce the content. As network service

is built out, the content delivery naturally increases at the same time.

Content-centric networking uses a practical data storage cache at each level of the network to dramatically decrease the transmission traffic, and also increase the speed of response. The cache envisioned by CCN is a packet-level cache present at each node in the tree of network equipment not a complete copy of some media file. In that way, the worst case is that everything behaves as it does now: A consumer requests some data and it propagates through the network. However, the second time the data is requested, if it is still in the cache at some level, there are dramatic savings. A number of studies have verified the potential benefits of this approach.

One important issue with content-centric delivery is to assure that the name of the content sufficiently describes the information. For example, the name should include a version, so data can be corrected. It should also include a cryptographic hash, so that the content can be authenticated. The name also needs detailed information about the decoding format. In the case of video, an especially demanding case, it should describe the transcoding format. The general scheme appears able to accommodate all three of the common video distribution systems: Individual transcoding streams, quality layers, and even dynamic recoding. In the case of dynamic recoding, perhaps only a single recoder is required, and then the network caches the new transcoding.

If the content is cryptographically signed, the consumer's equipment can also validate the data easily, and tell upstream devices when it is corrupted. Since the validation complaints can and should be verified at each node, the rejection will propagate only to correct caches with bad data. Hackers will not be able to deny service in the network with lies about hash validation.

A content-centric network is also well suited to provision mobile devices. When content is used, a connection does not need to follow the mobile device. The data needed can be requested anew, and may often be immediately available at new network nodes because it has already been cached.

In this model, the logical place to put commercial copy control and security is not in consumer equipment, but in the neighboring commercial network nodes. If the node agrees that the consumer has a distribution agreement, then restricted content can be delivered. Such delivery contracts require relatively few, cheap CPU cycles from devices already present near the edge of an ISP's net. If there are commercial restrictions, those may need to be included in the content names, as well.

See also [edit]

- Burnbit
- ed2k URI scheme
- Magnet URI scheme
- Metalink

Categories: Computer networking

This page was last modified on 18 April 2014 at 12:25.

Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.

Privacy policy About Wikipedia Disclaimers Contact Wikipedia Developers Mobile view